

Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/GB05/000726

International filing date: 28 February 2005 (28.02.2005)

Document type: Certified copy of priority document

Document details: Country/Office: GB
Number: 0404517.5
Filing date: 01 March 2004 (01.03.2004)

Date of receipt at the International Bureau: 18 April 2005 (18.04.2005)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)



World Intellectual Property Organization (WIPO) - Geneva, Switzerland
Organisation Mondiale de la Propriété Intellectuelle (OMPI) - Genève, Suisse



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.


In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 21 March 2005





Patents Form 1/77

THE PATENT OFFICE
PEThe
Patent
Office01MAR04 E877180-1 002776
P0171700 0.00-0404517.5 ALMOUNT CHAPatents Act 1977
Rule 16)

01 MAR 2004

RECEIVED BY FAX

Request for grant of a patent

(See the notes on the back of this form. You can also get
an explanatory leaflet from the Patent Office to help
you fill in this form)

1 MAR 2004

6404517.5 The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

IP/P7356

2. Patent application number

(The Patent Office will fill in this part)

0404517.5

3. Full name, address and postcode of the or of
each applicant (underline all surnames)

QINETIQ LIMITED

Registered Office 85 Buckingham Gate
London SW1E 6PD
United Kingdom

Patents ADP number (if you know it)

If the applicant is a corporate body, give the
country/state of its incorporation

GB

08178196002

4. Title of the invention

Threat mitigation in computer networks

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom
to which all correspondence should be sent
(including the postcode)

TOCHER ALASTAIR J et al

QINETIQ LIMITED

IP Formalities

A4 Bldg

Cody Technology Park

Ively Road

Farnborough

Hants GU14 0LX United Kingdom

Patents ADP number (if you know it)

08183873001

6. If you are declaring priority from one or more

earlier patent applications, give the country
and the date of filing of the or of each of these
earlier applications and (if you know it) the or
each application number

Country

Priority application number

Date of filing

(if you know it)

(day / month / year)

7. If this application is divided or otherwise
derived from an earlier UK application,
give the number and the filing date of
the earlier application

Number or earlier application

Date of filing

(day / month / year)

8. Is a statement of inventorship and of right
if to grant of a patent required in support of
this request? (Answer 'Yes' if:

a) any applicant named in part 3 is not an inventor, or

b) there is an inventor who is not named as an
applicant, or

c) any named applicant is a corporate body.

See note (d))

Yes (b)

Patents Form 1/77

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description 15

Claim(s) 02

Abstract 01

Drawing(s) 04 only *ll*

10. If you are also filing any of the following, state how many against each item.

Priority documents 0

Translations of priority documents 0

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 0

Request for preliminary examination and search (Patents Form 9/77) 0

Request for substantive examination (Patents Form 10/77) 0

Any other documents 0
(please specify)

11. I / We request the grant of a patent on the basis of this application

Signature

A J Tocher
A J Tocher

Date 01/03/2004

12. Name and daytime telephone number of person to contact in the United Kingdom
Mrs Linda Bruckshaw
01252 392722

Warning

After an application for a patent has been filed the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent of the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have attached 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Patents Form 1/77

P7356 - WYATT, G

THREAT MITIGATION IN COMPUTER NETWORKS

FIELD OF THE INVENTION

5 The present invention relates to apparatus, methods, signals, and programs for a computer for security threat mitigation and document transfer in computer and communications networks and systems incorporating the same.

BACKGROUND TO THE INVENTION

10 High resilience networks frequently have requirements for exchange of information with networks of low assurance, including networks of unknown threat level such as the public Internet. Traditionally, the approach to solving this problem is an air-gap between the two domains, with information exchanged between them on floppy disk. However this approach is both time-consuming and potentially risky.

15 This paper proposes alternative techniques to enable assured, two-way, information flow between high resilience networks and other networks of unknown threat. The techniques include conventional and novel technologies designed to control and constrain information formats, manage the environment between domains with network level controls, and provide assured, user-instigated, release sanctions.

20 Access to the public Internet for, for example, e-mail and web access is now almost essential even for security conscious end-users. There are other, innumerable, unclassified or low classification systems with which such security conscious networks must communicate, such as the news, weather, electronic library, route planning, and other public information networks.

25 There is therefore a requirement to transfer information between such low classification systems to higher classification systems that require higher assurance and a greater degree of resilience. In such a context, public Internet-connected systems are inevitably considered a high threat. Strict controls must be placed at the boundaries between these systems to prevent both the introduction of malicious content from the low system into the high system and the leakage of high data to the low system. Lower assurance arises due to the higher (or unknown) threat level, and in general is likely to lead to lower levels of resilience in applications.

30 Historically, the security separation problem has been solved by total electronic separation between the low and high networks. Such separation is sometimes referred to as an "air gap".

But experience has shown that a genuine air gap is not always practical, since the low classification information may have high value in the high system -- for example, weather data, news, collaborative planning information between organisations, and information from public agencies.

- 5 Figure 1 summarises information flow between such high domains 10 and low domains 20. High domains typically have clearly identified and limited users with shared motivation; the users' capabilities are known and controlled within the high domain. In contrast, low domains may have little control over who the domain users are, resulting in users with largely unknown and uncontrollable capabilities having access to the domain. It is typically such unidentifiable users who introduce the high threat level to such networks.

There is a recognised need then to provide a secure connection between such high and low networks in such a way that information can be exchanged without posing an unacceptable threat to the resilience of the high network.

- 15 It is well known to use floppy disks and CDs to transfer files between domains of unequal security level to bridge the air gap. However these simple techniques themselves introduce a number of risks and are inadequate to deal with modern threats which use multiple propagation mechanisms to gain access to networks.

Some of the dangers posed to a high assurance domain are reported widely and frequently. Examples include:

- 20
- attacks against Internet facing web servers to extract client credit card details;
 - e-mail arriving from the Internet carrying viruses that infect a business' Intranet;
 - the emergence of social engineering (sometimes referred to as 'phishing') attacks where the attacker masquerades as an on-line retailer or bank, sending e-mails to customers inviting them to click on a web hyper-link under the pretext of performing a necessary administrative task. The attacker, who owns the web site, can then harvest the usernames, passwords and credit card details of the victims and use these credentials to gain access to the high assurance system.
- 25

Attacks can be categorised in many ways, using terms such as:

- 30
- Viruses -- malicious code (sometimes referred to as "malware" which replicates itself to other host programs, areas of memory, disk boot sectors, or macro capable documents. Viruses may also execute a malicious payload.

- Worms – malicious code which makes copies of itself and can exploit program vulnerabilities to propagate. The propagation mechanism may alternatively be within the worm code itself.
- Trojan Horses – program which do not replicate themselves but can damage the host computer or use the host to launch further attacks, often under the direct control of the attacker.

Other kinds of attacks include back doors, rootkits, BIOS and Microcode malware, social engineering attacks, and buffer overflows.

Attacks reported in the press are generally the opportunistic type that become global phenomena, such as variants of the Sobig or LoveBug viruses. Targeted attacks are less widely reported and less well understood. It is therefore difficult to determine whether their apparent lack of frequency is due to their rarity, to the unwillingness of organisations to discuss such attacks publicly, or simply to their success (i.e. the attack succeeds and is so carefully concealed that it is never discovered).

The scattergun approach taken by opportunistic attackers is a time-consuming and troublesome nuisance to system administrators who must secure their Internet facing networks from the attacks. Where these networks are connected to an affiliated high domain with high-resilience requirements, the high domain administrators must, as a result, deal with the same threats posed by the opportunistic attackers, as well as addressing the potentially more devastating targeted attacks launched by skilled and motivated attackers.

It is well known that the majority of attacks on the Internet are from relatively unskilled attackers making use of publicly available tools and code that exploit known vulnerabilities, for which manufacturers' patches are generally available. However a highly motivated and highly skilled attacker could discover new vulnerabilities, develop new means to exploit those vulnerabilities, and attack a network using such novel methods which would avoid detection by existing commercial intrusion detection and filtering systems. Such commercial systems generally detect only known recognised patterns, or signatures, of attack.

A successful attack typically comprises one or more elements of the information security trinity of information confidentiality, integrity, and availability:

- Confidentiality – Information within the high domain should remain within the high domain unless its release to a lower domain is authorised and appropriate. For

example, personal information relating to the clients of a bank must remain confidential and not be leaked to the Internet, either accidentally or deliberately.

- 5 • Integrity – Information within the high domain should remain uncorrupted. An integrity attack might lead, for example, to a message 'Credit Joe Bloggs £3000.00' being changed to 'Debit Joe Bloggs £3000.00 resulting in inaccurate bank balance information being stored.
- 10 • Availability – Information services within the high domain must remain available. A well-known availability attack is the Distributed Denial of Service attack that has affected many Internet facing companies such as on-line banks and retailers. The servers of these companies are bombarded with bogus requests from thousands of computers infected with trojan horses controlled by the attacker. Valid user requests are unable to reach the server due to the overwhelming quantity of bogus traffic.

Some recent, successful and well-publicised attacks reveal a new trend for malicious code which:

- 15 • has combined characteristics of virus and worm for propagation, using mobile code (active content) that usually requires some form of user interaction, as well as an element involving network attack which can happen automatically;
- 20 • can embed itself into a system as a Trojan horse;
- can add a back door allowing a two-way communication channel back to 'base' (or more likely, a web site or IRC chat room in the attacker's control) where commands or updates to the virus code can be posted.

25 An interesting point about such code is the diverse range of attack vectors being used for propagation. Methods include direct communication using TCP/IP, application channels such as SMTP, and corrupted application data, all being used in combination. This kind of attack can propagate using floppy disks and CDs, therefore crossing what are perceived as "air gaps" so as to threaten the resilience of critical networks where users transfer files between high and low domains by those means.

30 The recent Sobig.F (August 2003) attack exhibited many such characteristics. It used several propagation techniques, including spreading via e-mail attachments and network shares, and it included an embedded SMTP engine. The payload was an URL downloader which, in effect, meant that the payload was infinitely variable, being dependent upon the imagination and talent

of the attacker to create their preferred attack and post the code to the web site from which the virus was pre-programmed to seek the download. The most obvious effect of the Sobig.F attack was denial of service through network flooding and e-mail system overload.

Furthermore, the number of malware attacks reported is increasing at an alarming rate. For example, the Computer Emergency Response Team (CERT) operated by Carnegie-Mellon University in the US reports the following figures:

- 21,756 viruses reported in 2000
- 114,855 viruses reported in 2003 (to October)

The propagation speed is also increasing as the number of attack vectors has increased, and as the time between the announcement of a vulnerability and the associated virus release has narrowed, sometimes to a few hours.

Yet, as noted above, such public statistics largely overlook the issue of targeted attacks, the prevalence and effects of which are almost completely unknown. Unfortunately, the threat of targeted attacks is the key concern to military networks.

SUMMARY OF THE INVENTION

According to a first aspect of the present invention there is provided a method for mitigating risks of connecting low and high assurance computer domains.

In particular, according to a first aspect of the present invention there is provided a computerised method of securely communicating an electronic document, the method comprising the steps of: where the data format of the document is not in a predetermined set of data formats, converting the document to a data format in the set; conveying the resulting document securely to a user sanction function for review and sanction by a user; if the document has been sanctioned, digitally signing the document for onward transmission to the recipient.

The sender and recipient may be in the same or different security environments which are subject to different levels of threat. The sender's security environment may be considered to be more 'secure' (i.e. having a lower threat) or less secure (i.e. having a higher threat) than the recipient's environment. The sender and recipient may also be in the same security environment. The invention described herein can be used in or for any such combination of environments. In one embodiment, the message is conveyed from an environment that is

considered to have a higher level of security to a recipient in an environment that is considered to have a lower level of security.

5 In an alternative embodiment, the message is conveyed from an environment that is considered to have a lower level of security to a recipient in an environment that is considered to have a higher level of security.

In one embodiment the method further comprises the step of: conveying the resulting document securely to a user sanction function for review and sanction by a user;

In a further embodiment, the document is transmitted via at least one one-way data diode.

10 In a further embodiment, the data is conveyed from the sender via electronic mail or by other file transfer mechanism.

In a further embodiment, the step of converting the document comprises performing a transformation which modifies the underlying data format of the document whilst substantially preserving the visible informational content.

15 In a further embodiment, the step of performing a transformation comprises applying an intentionally lossy compression method.

In a further embodiment, the step of adding entropy or randomness to certain elements of the document such as colour schemes, or spacing.

In a further embodiment, the method is performed in response to a user request to send a document by electronic mail to a recipient in a security domain distinct from that of the sender.

20 According to a second aspect of the present invention there is provided a computerised method of providing security in a computer network, the method comprising the steps of: where the data format of a document for transmission between parties over the network is not in a predetermined set of data formats, converting the document to a data format in the set; conveying the resulting document securely to a user sanction function for review and sanction
25 by a user; if the document is sanctioned, digitally signing the document for onward transmission to the recipient.

The invention also provides for systems for the purposes of communications and which comprise one or more instances of apparatus embodying the present invention, optionally combined with other additional apparatus.

In particular, according to a third aspect of the present invention there is provided a computer system arranged to perform the preceding methods.

The invention also provides for computer software in a machine-readable form and arranged, in operation, to carry out each function of the apparatus and/or methods. In this context this
5 includes not only source and object code, on a machine readable medium, for execution on a general purpose computer but also code intended for computer simulation of such systems or for compilation to silicon whereby the program may be implemented as a computer chipset (including the chipset comprising a single chip).

In particular, according to a fourth aspect of the present invention there is provided a computer
10 chipset arranged to perform the preceding methods.

In particular, according to a fifth aspect of the present invention there is provided a program for a computer arranged to perform the preceding methods.

The invention is also directed to novel signals employed by the other aspects of the invention.

The preferred features may be combined as appropriate, as would be apparent to a skilled
15 person, and may be combined with any of the aspects of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

In order to show how the invention may be carried into effect, embodiments of the invention are now described below by way of example only and with reference to the accompanying figures in which:

20 Figure 1 shows a schematic diagram of information transfer between domains;

Figure 2 shows a schematic diagram of a first system in accordance with the present invention;

Figure 3 shows a schematic diagram of a second system in accordance with the present invention;

25 Figure 4 shows a schematic diagram of a second system in accordance with the present invention.

DETAILED DESCRIPTION OF INVENTION

The present inventors have identified three categories of techniques that can be used to control information flow between domains of differing security levels:

- Data format control
- Environment control
- User control and release sanctions

Data format control techniques involve three related processes:

- 5 • using inherently 'safe' formats,
- format conversion by which data is transformed from one format to another format which utilises a different format 'grammar' and 'syntax' from that used by the original format, and
- checking that the formatting rules have been obeyed.

- 10 The dangers of some complex data formats are known. For example, HTML used in web pages and e-mail can contain active content described with a scripting language which, by default, is processed and rendered by the client machine. Many attacks exploit vulnerabilities in the script interpreter or the host application to run malicious and damaging program code. Similarly, complex formats used to describe documents created with a word processor can
- 15 contain macros which again perform malicious actions. Word processor formats also allow for 'hidden content', where multiple versions of a document are embedded within a single document description. Such files may contain data previously deleted by the user and not visible on the screen, but still encoded within the file.

Formats that ban or tightly constrain active content are safer and are in general to be preferred.

- 20 Examples of such formats include earlier versions of PDF (PostScript Distribution Format), which tightly constrained the use of scripting, and banned dangerous functions such as general file access. However, such formats cannot be regarded as entirely safe since any complex language allows for the possibility of malformed data structures; these can produce unexpected behaviour in end systems, such as denial of service or buffer overflow attacks, which in effect
- 25 convert a passive data structure into active code. PDF is known to suffer from such problems, and there have been documented examples of such abuse. Complex protocol specification languages such as ASN.1 (Abstract Syntax Notation One) have also been shown to suffer similar weaknesses.

- 30 One approach to solving these problems is to write document viewers and protocol implementations that are carefully crafted and exhaustively tested. In general, industry is increasingly recognising the need for this, but complex languages cause the data-space

requiring testing to be huge, or in some cases unbounded. Therefore, the assurance of such applications is limited, and something further is required in high assurance environments. A 'safe', or at least safer, format should be incapable of such (albeit unintended) subterfuge, using simple data structures with a well-defined syntax and grammar with a finite space.

- 5 Examples of simple formats include ASCII text and (certain kinds of) bitmap images. Their simplicity aids another part of format control, the format checking.

The idea behind format conversion is that by transforming the representation of data from one form to another, possibly with several iterations, any malware capabilities contained in the original information format would be lost, especially if one of the transformation processes
10 introduces an unpredictable, random element to the conversion process.

There are two stages at which format conversion might be used in a high assurance scenario.

- First, many complex formats can be converted into ASCII text or bitmap images. For example in a Bitmap image of a Word document, all the user-visible informational content remains (the document can still be read), but all the hidden content encoded in the complex Word format is
15 removed, along with the associated threat posed by that hidden content.

There are a number of commercial programs which can be used to take bitmap images of complex documents (in HTML, Word, Excel, Adobe Acrobat, and other formats) in a more sophisticated manner than a simple 'Print Screen' request. For example, these 'screen-scrape' programs can open a Word document so that the entire document is converted to a Bitmap
20 image, not just the single page visible on the screen.

In addition, some of the screen-scrape programs incorporate a 'text-scrape' function where the text within a Word or Adobe™ formatted document can be converted to ASCII text, again removing any hidden content, though at this stage any text-scraped scripting code, for example, may be preserved.

- 25 It is worth differentiating between the effect of a screen scrape and a text scrape. The text scrape takes a document and translates the complex encoding used by the word processor to represent letters of the alphabet and translates that into simple ASCII encoding. However both encodings are used to represent letters of the alphabet. Converting a word-processed document into a (bitmap) image profoundly alters the encoding; whilst the word processor
30 encodes letters of the alphabet, the image encodes only the colour and brightness of pixels in the image. It is only in the viewer's brain that those encodings are translated back into text.

The second time a format conversion might be performed is when releasing bitmap images, just after the format checking is completed. At this stage, the bitmap image may be transformed into a JPEG image. The transformation process, using a lossy compression algorithm like JPEG, would mean that while the information content would remain the same (the JPEG and bitmap images would appear almost identical to the viewer) the data format would undergo a transformation.

This second transformation makes it even less likely that an opportunistic attacker's hidden or malicious content can be preserved, particularly given that it may be irretrievably lost thanks to the lossy compression. Compression is of course doubly desirable given that bandwidth restrictions would make the exchange of large bitmap files difficult to scale to a large user community.

To defeat the targeted attacker it would also be possible to introduce some random alterations to the bitmap image prior to the (optional) JPEG conversion; for example, by making small, random modifications to the brightness values of some, or all, of the pixels in the bitmap, and even to the image dimensions by adding a randomly sized border to the image. The changes would be visually indiscernible to the viewer of the final image, but make it very difficult for an expert attacker to predict the exact output of the conversion and so design an attack to exploit a hypothetical vulnerability in the recipient's image viewer. Indeed the variation introduced may be varied pseudo-randomly over time to further confound attackers.

Crucially, the initial screen or text scrape conversion process does *not* need to be trusted or assured. However the optional process to randomise Bitmaps and the format checking stage is preferably assured.

Regarding format checking, there are an increasing number of products which are designed to check complex formats such as HTML, XML, Word, or Excel for hidden content, malicious macros, and the like. Such checkers are extremely useful and valuable in scenarios where this richness of information transfer is permissible. They can use simple heuristics such as ensuring that Word documents have not been "fast saved". But simple heuristics are inadequate to address the full range of kind of threats discussed above.

For these situations, much more constrained data formats are required. If, for example, only two simple information formats such as ASCII text and bitmap images are allowed to be exchanged between the two domains, the scale of the problem is much reduced.

By only permitting ASCII text messages and the simplest variety of bitmap images to pass between domains, there is a significant reduction in the complexity required of the content

checkers. For example, an ASCII text checker may comprise one very simple function that would eliminate from the message any characters that were not from a recognised alphanumeric list or from a small subset of permitted punctuation markings. The punctuation could even be replaced as part of the conversion by words, as in old-fashioned telegrams ("STOP" for
5 ".", etc.). The presence of other content raises a warning and is removed from the message before being passed on. Therefore if for example the message contains Javascript, the opening and closing brackets ('<' '>') used to denote a 'tag' may be removed, rendering the Javascript incomprehensible to the receiving application. An optional, simple ASCII text search warns of the presence of blacklisted words.

- 10 A bitmap checker may be simpler still, provided the simplest kind of raw bitmap format is used exclusively. Such simple bitmaps are a rudimentary image format to parse. They contain a definition of the dimension of the image (numbers of pixels across and high) followed by the colour parameters for each of the pixels (defined for example in terms of the strength, from 0 to 255, of the Red, Green and Blue constituent colours). The bitmap format checker performs a
15 simple grammar and syntax check of each image so that each file is known to conform to these bitmap encoding rules.

Regarding environment control, the techniques of format control discussed above must be carefully managed and protected to ensure that they cannot be subverted. Here we specifically refer to the environment at the interfaces between the high and low assurance domains.

- 20 Traditional techniques to control the interface environment include firewalls, one-way data diodes (where information can only flow in one direction, for example from a low to high domain or vice versa), and the use of De-Militarised Zones, generally bounded by two firewalls or data diodes (or both).

- Such a controlled environment, may be constructed using commercial off-the-shelf (COTS)
25 components, some of which have formal (EAL) assurance rating. Such components include QinetiQ's one-way data diode (SyBard::Suite@ Diode) and the SWIPSY firewall toolkit. The SWIPSY toolkit is an E3 (equivalent to EAL4) evaluated product which allows additional code to be added to its security compartments without affecting the evaluation status of the toolkit itself. SWIPSY has security properties which assure network and process separation: processes
30 communicating with one network (for example the high domain) cannot communicate directly with the other network (for example the low domain) other than by via trusted mover agents which in turn force data to be passed to the format and content checkers.

Environment controls however are insufficient in isolation to control some malware, since the controls typically defend at the network level. However, when combined with application level

'checker' software and novel techniques described herein, they are vital in forcing the data through the checking processes. A number of other mechanisms must also be used if high assurance is required. These include assured user sanction, and potentially, machine virtualisation and these are discussed in more detail below.

- 5 Regarding user control and release sanctions, the assured intervention of a human user is critical for the release of data from a high domain to a low domain for two reasons. Firstly, if every exchange of information from high to low is governed by user sanction, the process can be meaningfully audited, making the users accountable for their actions. Secondly, a properly implemented and assured user sanction mechanism prevents any high domain machine
- 10 infected by a back-door attack from communicating with the low domain by any route other than via the user sanctioned channel. With this restriction on information flow in place, as well as the format checking and format conversion, it becomes a significantly more challenging task for a back door program to communicate with the low domain attacker without attracting the attention of the sanctioning user.
- 15 The software used to implement the user sanction process should be simple and trusted. Firstly, the data should be presented to the sanctioning user for approval and release. This data should be confined to simple formats such as ASCII text or bitmap images as described above. Secondly, the user could, optionally, instruct the software to digitally sign the data to prevent modifications between the user's desktop and the domain boundary. The user must be
- 20 able to view everything on the screen that is to be signed. This implies that a trusted viewer is used: that is, an assured computer that the user can trust will display the complete content of the message and neither this viewer nor the signer can be subverted by any attack.

The idea of a trusted signing device is a recognised requirement in some civil applications. The EU Digital Signature Directive recognises the need for a Secure Signature Creation Device

25 (SSCD), and some EU member states have enshrined this requirement in their legislation. In addition, industry seems to be following a similar trend: the Trusted Computing Group's (TCG) Trusted Platform Module (TPM) specification includes "Data Attestation" facilities for signing data structures; and recently published documents describing Microsoft's Next Generation Computing Base (formerly Palladium™) detail aspirations for both trusted viewers and trusted

30 signing mechanisms, as part of the trusted part of the operating system.

Referring now to Figure 2, a first embodiment of an architecture for threat mitigation in computer networks – between a high-domain 10 comprising at least one client machine 11 and a low-domain 20 comprising at least one client machine 22 – comprises a format converter 12 and a user sanction and release function 13.

To send data from the high domain to the low domain, an end user connects to a high domain web server 12 that can perform format conversion. Format conversion may be in the form of screen-scraping or text-scraping, essentially converting any complex file format (e.g. a Microsoft Word™ document) into a reduced functionality document (e.g. a bitmap image or ASCII text document). The converted file is then routed, via e-mail, to the user sanction release mechanism 13. The trusted user sanction device may, for example, be built upon the Trusted Solaris™ operating system (EAL4 rating) and may include an associated firewall function 40 (e.g. built upon the EAL4-rated SWIPSY toolkit).

Upon this system a console-based viewing program installed. The console viewer strips any MIME (or similar) encoding surrounding the message body (or attachment in the case of a bitmap image), performs the format check, and presents the ASCII text or bitmap to the user.

The reviewing user – who can be the original sender or another independent individual – can then choose to stop or release the message. If the release option is chosen, the message is digitally signed and packaged in an S/MIME envelope.

The digital signature is applied 31 using an assured cryptographic library (for example the CESC Cryptserve algorithm suite - assured to EAL4) and suitable toolkit to create the email (e.g. S/MIME) envelope. In one specific embodiment a copy of the "To:", "From:" and other header fields from the message body are included in the envelope in order to counteract the known S/MIME weakness where header information is unsigned. However, this does not affect correct reception of the S/MIME message and checking in a standard S/MIME client.

The signed message is then forwarded to the DMZ 30 and to another SWIPSY firewall 50 (this one operating without human intervention). At the DMZ firewall the digital signature is verified 31 and (optionally) removed before it is forwarded to the low domain SMTP server.

Referring now to Figure 4, the communication route from a low domain to a high domain does not necessarily require a user sanction, although there may be good security reasons why such intervention may be required in some circumstances. This process can therefore be fully automated.

A low-domain user connects to a low domain web server that can perform screen-scrape or text-scrape 22 to converting a complex file format (e.g. a Microsoft Word document) into a simple format such as a bitmap image or ASCII text document.

The scraped file is sent through a one-way data diode 60 to the trusted format checker 32 hosted by a SWIPSY machine in the DMZ 30. The checker performs a content check on the ASCII or bitmap file as before.

5 If the content check is successful, the bitmap file may be 'randomised' and optionally transformed into another image format (e.g. JPEG format). The process of randomising introduces some variations in the original encoding to mitigate any residual risk of threats being conveyed in the precise detail of the file layout.

10 The transport protocol encodings that surrounded the file (e.g. MIME encoding and SMTP headers) may then be stripped off and recreated on the SWIPSY machine 32. Alternatively, protocols (such as FTP) which require no transport 'envelope' around the file can be used.

The checked file may then be passed to the high domain 10 through a second one-way data diode 70.

The file is delivered to the high domain server from which the client can collect the 'safe' file.

15 In an optional additional step, the file may be transformed into the original, or other higher-functionality document format (e.g. Microsoft Word format) at this stage. This might be achieved using, for example, known Optical Character Reader (OCR) software to recreate an editable file from a bitmap image. This could happen in an automated fashion in some point in the destination domain, whether upon entry to the destination domain, upon delivery to the recipient or any intervening point. This may in some circumstances reduce the robustness of
20 the architecture without this step since it might prove possible for any 'malware' in the original document to be reconstructed as part of this process.

25 It can be seen that the data diodes 40, 50, 60, 70 illustrated in the diagrams above are directed in opposite directions, thereby creating a two-way flow of data between high and low that might appear to make the data diodes redundant. However, the high and low domains illustrated in the diagrams are generic labels and not meant to imply that the low domain is the same low domain in both diagrams. Rather, the scenarios described in this paper are assuming the high domain has a multiplicity of logically distinct connections to lower domains.

30 In the embodiments described above an architecture is presented which imposes stringent constraints on the information transfer between high and low domains. Other embodiments may require a 'richer' exchange of information, involving more complex file types and data encodings than simple ASCII text and Bitmap images. Examples of such formats include PDF files and database updates. For such kinds of files, complex file types may be transformed into

a single, standard XML representation. Such a complex – but uniform – XML would then be subjected to content checking on a trusted platform, and then transformed from the XML back to the original encoding (e.g. Microsoft Word) or to a new encoding (e.g. PDF format).

- Machine Virtualisation allows a second complete operating system to be installed on a user's machine and for that second operating system to run concurrently as a second 'virtual' machine with the host operating system. Virtual machine software such as VMWare™ allows many such virtual machines to operate simultaneously. Virtual machine technology may be used to create secure multi-system desktops where the two virtual machines are separated by known assured mechanisms.
- 10 Preferably, the trusted user sanction mechanism is built into a separate virtual machine on the client's desktop. Such an arrangement may make use of an assured and highly locked-down Operating System that is specifically for this one purpose. Such an architecture adds flexibility and convenience, allowing users to release documents from their own desktops, either to replace or supplement the independent trusted 'domain' signatory described above. Where it
- 15 supplements the independent signatory it effectively enforces a two-person rule release mechanism.

In summary then, a combination of techniques is proposed to allow a potentially assured, two-way exchange of information between high and low assurance domains. The architecture potentially offers greater security than an air-gap since whilst an air-gap is an attractive idea in theory it is often difficult to achieve in practice: pragmatic users faced with an air-gap between security domains will often reach for a floppy disk or USB memory stick rather than re-type the entire document from the high machine to the low.

20

The combination of format controls, environment controls, and user controlled release sanctions outlined above have been designed to offer a secure solution to the urgent and unavoidable business requirement to share information.

25

Any range or device value given herein may be extended or altered without losing the effect sought, as will be apparent to the skilled person for an understanding of the teachings herein.

CLAIMS

1. A computerised method of securely communicating an electronic document, the method comprising the steps of:

where the data format of the document is not in a predetermined set of data formats,
5 converting the document to a data format in the set;

conveying the resulting document securely to a user sanction function for review and sanction by a user;

if the document is sanctioned, digitally signing the document for onward transmission to the recipient.

10 2. A method according to claim 1 in which the message is conveyed from an environment that is considered to have a higher level of security to a recipient in an environment that is considered to have a lower level of security.

3. A method according to claim 1 in which the message is conveyed from an environment that is considered to have a lower level of security to a recipient in an environment that is
15 considered to have a higher level of security.

4. A method according to any preceding claim further comprising the step of:

conveying the resulting document securely to a user sanction function for review and sanction by a user;

5. A method according to any preceding claim further comprising the step of:

20 transmitting the document via at least one one-way data diode.

6. A method according to any preceding claim in which the data is conveyed from the sender via electronic mail, or other file transport mechanism.

7. A method according to any preceding claim in which the step of converting the document comprises performing a transformation which modifies the underlying data format of the
25 document whilst substantially preserving the visible informational content.

8. A method according to any preceding claim in which the step of converting the document comprises adding entropy or randomness to certain elements of the document, such as colour schemes, or spacing.

9. A method according to claim 4 in which the step of performing a transformation comprises applying an intentionally lossy compression method.

10. A method according to any preceding claim performed in response to a user request to send a document by electronic mail to a recipient in a security domain distinct from that of the sender.

11. A computerised method of providing security in a computer network, the method comprising the steps of:

where the data format of a document for transmission between parties over the network is not in a predetermined set of data formats, converting the document to a data format in the set;

conveying the resulting document securely to a user sanction function for review and sanction by a user;

if the document is sanctioned, digitally signing the document for onward transmission to the recipient.

12. A computerised method of providing security in a computer network, substantially as described in the foregoing text and/or with reference to the accompanying figures.

13. A computer system arranged to perform the method of any one of claims 1- 12.

14. A program for a computer arranged to perform the method of any one of claims 1-12.

15. A computer chipset arranged to perform the method of any one of claims 1-12.

ABSTRACT


Figure 2 to accompany Abstract

A computerised method – and related apparatus, systems, programs for a computer and signals – for securely communicating an electronic document between high and low security domains.

- 5 Where the data format of the document is not in a predetermined set of simple data formats (for example bitmap images or ASCII text), the document is automatically converted to a data format in the set. Optionally a "lossy" transformation may also be applied to further confound attackers. The document is then conveyed securely to a user sanction function for review and sanction by a human user. Once sanctioned the document is digitally signed for onward
- 10 transmission to the recipient. For transmission from low domain to high domain, user sanctioning of document release may be omitted.

1/4

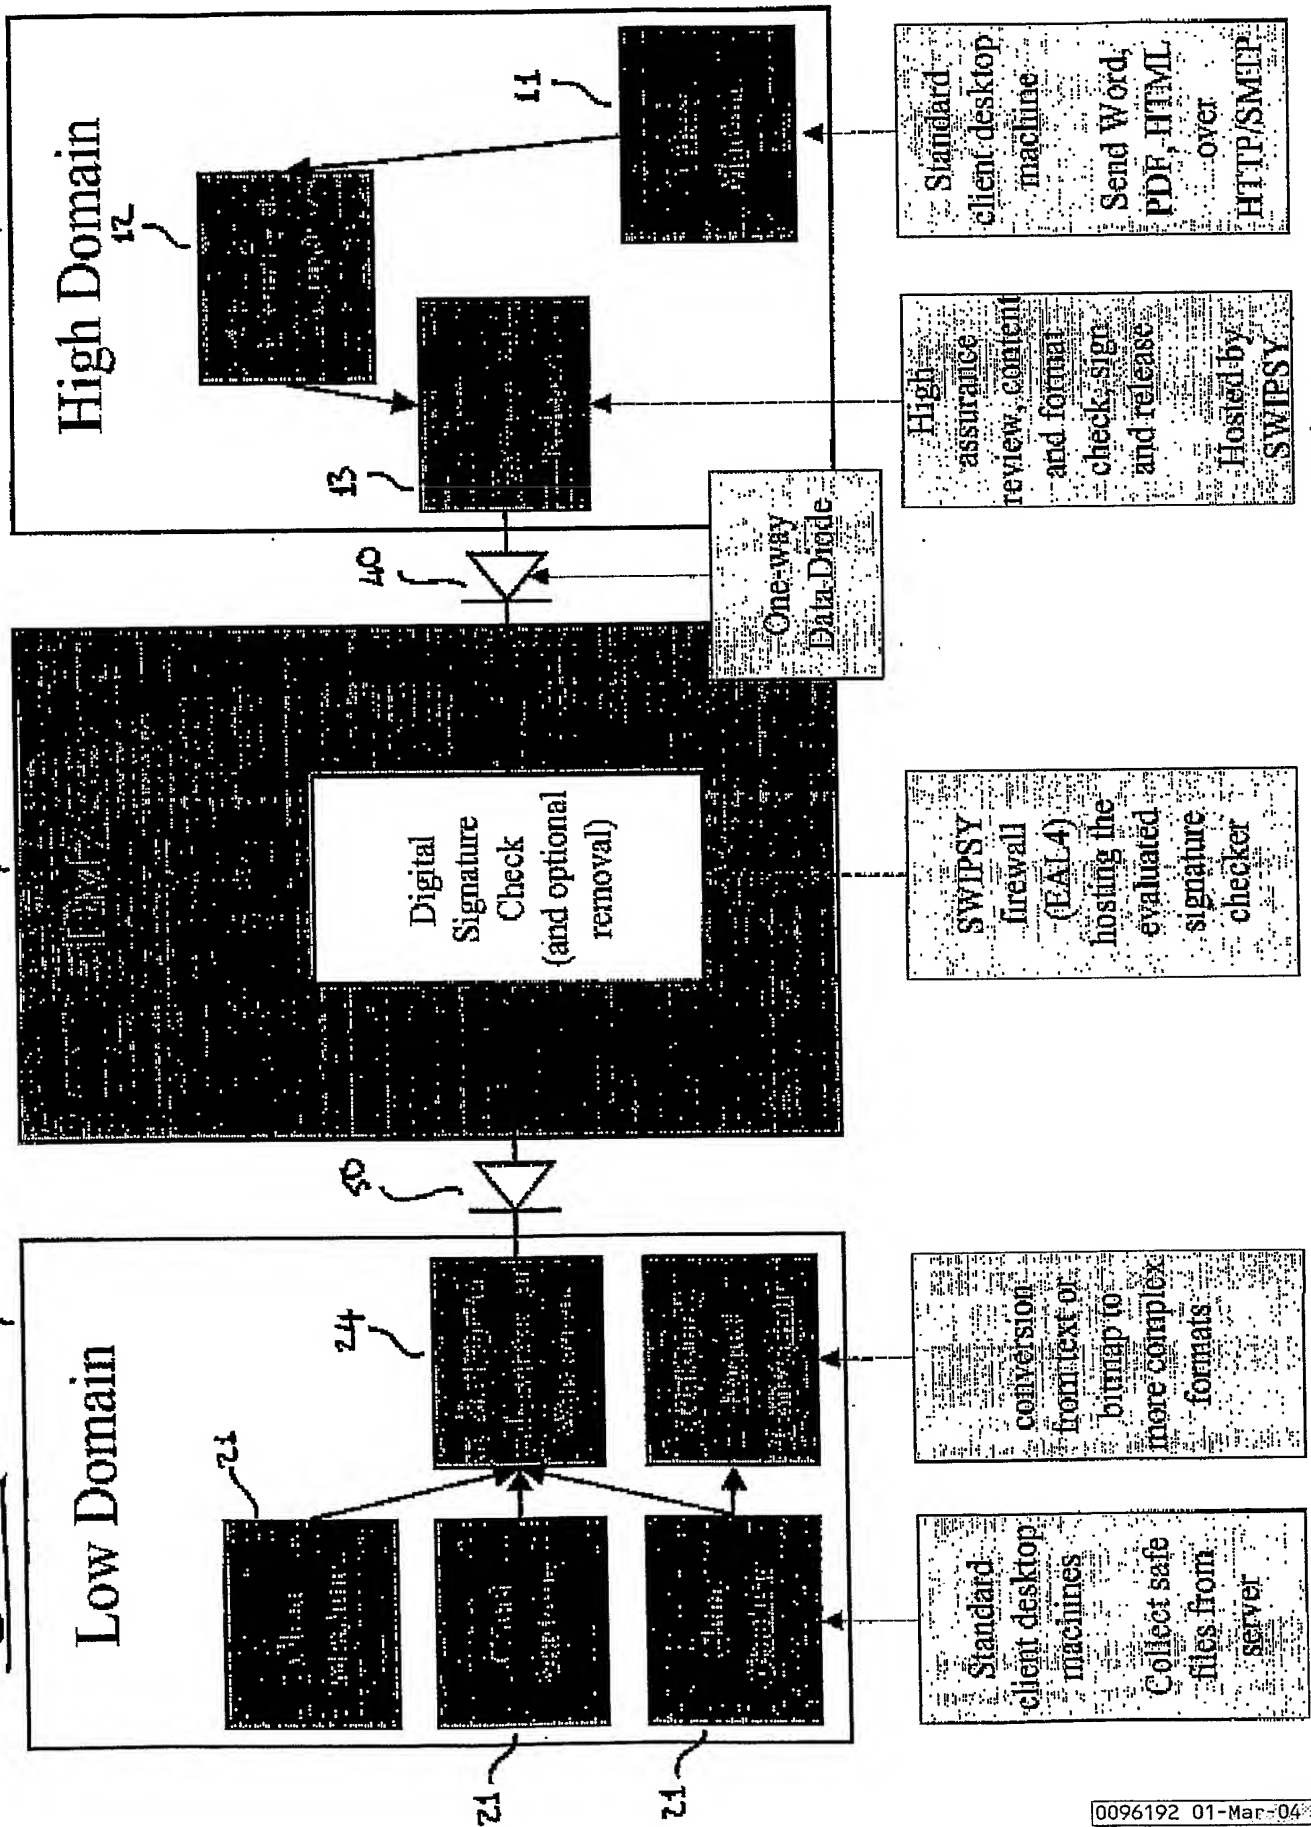
Figure 1





2/4

Figure 2





3/4

Figure 3

13

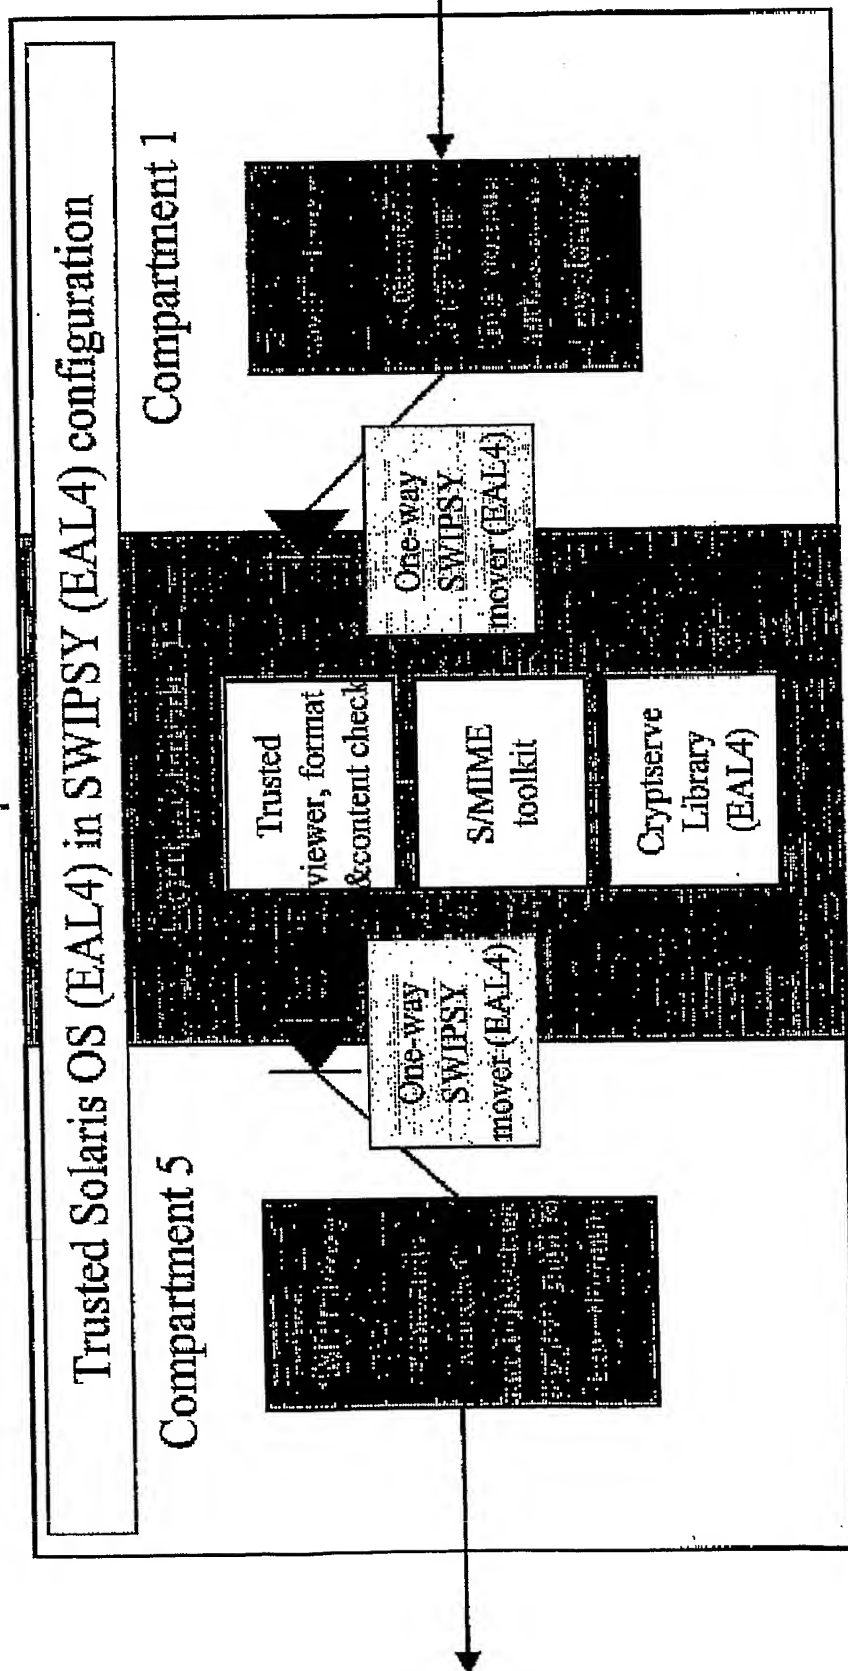




Figure 4

20

30

10

